M^{2}S-Net: Multi-Modal Similarity Metric Learning based Deep Convolutional Network for Answer Selection

Lingxun Meng[1] **Yan Li**[2]
[1]Business Advertising Part, Sogou Network Technologies,
Sogou-Inc, Beijing, 100084, China.
menglingxun@sogou-inc.com
[2]Key Lab of Intelligent Information Processing of Chinese Academy of Sciences (CAS),
Institute of Computing Technology, CAS, Beijing, 100190, China.
yan.li@vipl.ict.ac.cn

arXiv:1604.05519v2 [cs.CL] 24 Apr 2016

Abstract

Recent works using artificial neural networks based on distributed word representation greatly boost performance on various natural language processing tasks, especially the answer selection problem. Nevertheless, most of the previous works used deep learning methods (like LSTM-RNN, CNN, etc.) only to capture semantic representation of each sentence separately, without considering the interdependence between each other. In this paper, we propose a novel end-to-end learning framework which constitutes deep convolutional neural network based on multi-modal similarity metric learning (M^{2}S-Net) on pairwise tokens. The proposed model demonstrates its performance by surpassing previous state-of-the-art systems on the answer selection benchmark, i.e., TREC-QA dataset, in both MAP and MRR metrics.

1 Introduction

Inspired by the achievements of convolutional networks (a.k.a, ConvNets) in the field of computer vision, more and more researchers constitute ConvNets for kinds of natural language processing tasks, e.g., text classification (Zeng et al., 2014; Kim, 2014), text regression (Bitvai and Cohn, 2015), short text pair re-ranking (Yih et al., 2013; Severyn and Moschitti, 2015), and semantic matching (Hu et al., 2014).

For the answer selection task, i.e., given a question and a set of candidate sentences, choosing the correct sentence that contains the exact answer and can sufficiently support the answer choice, most of the previous methods constitute Siamese-like architectures (like LSTM-RNN, CNN, etc.) to learn the semantic representation for each sentence, and then use cosine similarity or weight matrix to compute the similarity of the pairwise representations (Feng et al., 2015; Severyn and Moschitti, 2015; Wang and Nyberg, 2015). At the same time, these works mostly adopted shallow architectures for sentence modeling, since deeper nets did not bring better performance. On the contrary, we firmly convinced that one can benefit more from deep learning strategy.

Following the success of RNN-based attentive mechanism designed for machine translation task (Bahdanau et al., 2014), recently some works attempted two-way attention mechanism for sentence pair matching problems (Tan et al., 2015; Santos et al., 2016; Yin et al., 2015). Specifically, these works mainly utilize the information from one input to directly inference the computation of the other's representation, and the attention vectors are computed along with the distances between the segments in both directions, and also frequently along with the layer-wise representation construction. As a result, such soft attention mechanism boosts a great improvement in performance, yet aggravates computations and model complexity.

Since the interdependence within the sentence pair has proven to be useful, why not we directly model the matching strategy, i.e., constitute an end-to-end deep network to simultaneously model the semantic representation and similarity metric? In this paper, we propose a novel learning framework for answer selection task. Different from previous popular networks for sentence semantic learning, we constitute a deep convolutional neural network based on pairwise token matching with multi-modal similarity metric, named by us M^{2}S-Net, where the learnt multi-modal similarity metric provides a comprehensive and multi-granularity measurement. Experimental results on the benchmark dataset of answer selection task indicate that the proposed model can greatly benefit from deep network structure as well as multi-modal similarity metric learning, and also demon-

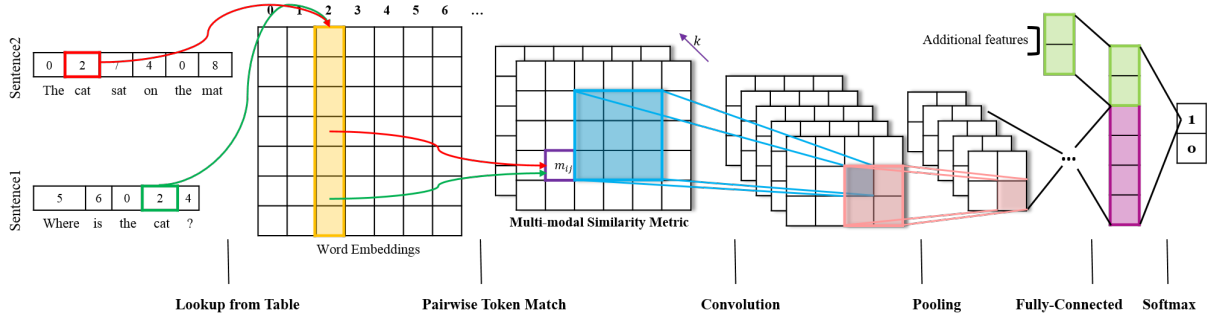

Figure 1: Our M^{2}S-Net for sentence pair matching.

strate that the proposed M^{2}S-Net outperforms a variety of powerful baselines and achieve state-of-the-art.

2 M^{2}S-Net

Siamese-like network is a classic structure for sentence pair matching. Typically, one convolutional and pooling layer or recurrent layer is used to model the semantic learning for each sentence, and pair similarity is calculated with cosine similarity based on the learned semantic representations. Considering that attention mechanism has been introduced successfully to positively influence each representation learning by computing the interdependence with the other one, in this paper, we propose a novel learning framework for sentence pair matching, where the pairwise token similarity is computed to generate similarity matrix firstly, and then a deep convolutional network is constructed to learn pairwise matching representation, finally concatenate the learned pairwise matching representation and additional simple word-level overlap feature to feed into a *pointwise* rank loss function (please see Fig. 1 for better understanding).

2.1 Pairwise Token Matching

As a fundamental component in M^{2}S-Net, it is of vital importance to design an appropriate similarity measurement f_{match} for pairwise token matching. Specifically, we design three kinds of match-score, i.e., Euclidean similarity, cosine similarity, and learnt similarity metric.

For a sentence pair $S_1 = \{\boldsymbol{w}_1^i, i \in [0, l_1 - 1]\}$ and $S_2 = \{\boldsymbol{w}_2^j, j \in [0, l_2 - 1]\}$, where l_1, l_2 are the word length of sentence S_1, S_2, respectively, $\boldsymbol{w}_1^i, \boldsymbol{w}_2^j$ are d-dimensional word embeddings W pre-trained under vocabulary V, we first calculate the similarity matrix $M \in \mathbb{R}^{l_1 \times l_2}$, where each element $m_{i,j} \in M$ is computed according to the three measurements we adopt, as follows:

$$m_{i,j} = \begin{cases} 1/(1 + |\boldsymbol{w}_i^1 - \boldsymbol{w}_j^2|) & \text{(Euc.)} \\ (\boldsymbol{w}_i^{1T} \boldsymbol{w}_j^2)/(||\boldsymbol{w}_i^1|| \cdot ||\boldsymbol{w}_j^2||) & \text{(Cos.)} \\ \boldsymbol{w}_i^{1T} U \boldsymbol{w}_j^2 + b_{i,j} & \text{(Metric.)} \end{cases} \quad (1)$$

It should be noticed that multi-modal metric could be used here for enriching the matching representation, i.e., the matrix $U \in \mathbb{R}^{k \times d \times d}$, and the corresponding bias term $B^{(m)} \in \mathbb{R}^{k \times l_1 \times l_2}$, where k should be assigned as the amount of similarity metric (the Euclidean and cosine measurements are the special cases where $k = 1$). Thus, the richer matching matrix comes in the form of $M \in \mathbb{R}^{k \times l_1 \times l_2}$, and going to be fed into the following convolution layer.

2.2 Convolution and Pooling

The convolution layer in this work consists of a filter bank $F \in \mathbb{R}^{n \times c \times h \times w}$, along with filter biases $\boldsymbol{b} \in \mathbb{R}^n$, where n, w and h refer to the number, width and height of filters respectively, and c denotes the channels of data from the lower layer. More specifically, for the first convolution layer, c equals to the multi-modal parameter k, which means convolving across all the similarity modals to learn the pattern. Given the output $L^{t-1} \in \mathbb{R}^{c \times l_h \times l_w}$ (L^0 represents similarity matrix M) from the lower layer, the output of the convolution with filter bank F is computed as follow:

$$\begin{aligned} L^t &= \tanh(F * L^{t-1} + \boldsymbol{b}) \\ &= \tanh([\boldsymbol{f}_i^T \boldsymbol{l}^{t-1}_{c \times (j-h+1:j) \times (l-w+1:l)} + b_i]) \end{aligned} \quad (2)$$

where * is marked as the convolutional operation, i indexes the number of filters, j and l indicates the sliding operations for dot production along the axis of width and height with one step size.

Typically, there exist two types of convolution: *wide* and *narrow*. Even though previous works (Kalchbrenner et al., 2014; Denil et al., 2014) have pointed out that using *wide* type of convolution was able to better and more frequently reach boundaries of sentences than the *narrow* type, and could ignore the requisition of the *narrow* type that filter width must be smaller than input data width (i.e., $l_h > h$ and $l_w > w$) to guarantee a valid non-empty result, we use the *narrow* type for convenience. Finally, we get the output of layer t as $L^t \in \mathbb{R}^{n \times (l_h - h + 1) \times (l_w - w + 1)}$.

The outputs from the convolutional layer (passed through the activation function) are then fed into the pooling layer, whose goal is to aggregate the information and reduce the representation. Technically, there exist two types of pooling strategy, i.e., *average* pooling and *max* pooling, and both pooling methods exhibit certain disadvantages: in *average* pooling, all elements of the input are considered, which may weaken strong activation values; the *max* pooling is used more widely and does not suffer from the drawbacks of average pooling. However, as shown in (Zeiler and Fergus, 2013), it can lead to strong overfitting on the training data and, hence, poor generalization on the test data. For simplification and stability, we adopt the *average* pooling strategy in this work.

2.3 Batch Normalization

Batch Normalization (BN) (Ioffe and Szegedy, 2015) was originally proposed to reduce the changes in distribution of each layers input during training, which was called *internal covariance shift* . As an important component of deep networks, it allows us to use much higher learning rates and be less careful about initial data preprocessing and weights initialization. Furthermore, it can also act as a regularization, in some cases eliminating the use of dropout (Srivastava et al., 2014).

Specifically, given the output of a convolutional layer $L \in \mathbb{R}^{c \times l_h \times l_w}$, we will normalize each channel as follow:

$$\hat{L} = [\hat{\boldsymbol{l}_i}] = [\frac{\boldsymbol{l}_i - \mathrm{E}[\boldsymbol{l}_i]}{\sqrt{\mathrm{Var}[\boldsymbol{l}_i]}}], \quad (3)$$

where i indexes the channel and $\boldsymbol{l}_i \in \mathbb{R}^{(l_h * l_w)}$ is the vectorization of each feature map, the expectation and variance are computed over each channel.

Set	#Question	#QApairs	%Correct	Judge
Train-All	1,229	53,417	12.0%	auto
Train	94	4,718	7.4%	man
Dev	82	1,148	19.3%	man
Test	100	1,517	18.7%	man

Table 1: Statistics of the answer sentence selection dataset. Judge denotes whether correctness was determined automatically (auto) or by human annotators (man).

A pair of parameters $\boldsymbol{\gamma}_i$ and $\boldsymbol{\beta}_i$ (with the same dimension with $\hat{\boldsymbol{l}_i}$) are further used to scale and shift the normalized value as follow:

$$\mathrm{BN}(\hat{L}) = [\boldsymbol{\gamma}_i \cdot \hat{\boldsymbol{l}}_i + \boldsymbol{\beta}_i]. \quad (4)$$

Above parameters are learnt along with the whole model training, and it usually suggests that Batch Normalization should better be used before nonlinearity activation and covers both fully connected layers and convolutional layers.

2.4 *Pointwise* Learning to Rank

We adopt simple *pointwise* method to model our answer selection task, though *pairwise* and *listwise* approaches claim to yield better performance. We deploy the cross-entropy cost to discriminatively train our framework as follow:

$$C_\theta = -\frac{1}{N}\sum_{i=1}^{N}[y_i \log p_i + (1 - y_i)\log(1 - p_i)] + \lambda(\|U\|^2 + \|B\|^2) \quad (5)$$

where p_i is the output probability of i^{th} sample through our networks, y_i is the corresponding ground truth, and θ contains all the parameters optimized by the network, i.e., $\theta = \{W; U; B; [F]; [\boldsymbol{b}]; [\boldsymbol{\gamma}]; [\boldsymbol{\beta}]\}$, λ is set to be $5e^{-4}$.

We use Stochastic Gradient Descent (SGD) to optimize our network, and AdaDelta (Zeiler, 2012) is used to automatically adapt the learning rate during the training procedure. For higher performance, hyper-parameter selection is conducted on the development set, and BN layer after each convolution layer is also added to speed up the network optimization. In addition, dropout is applied after the first hidden layer for regularization, and early stopping is used to prevent over fitting with a patience of 5 epochs.

3 Experiments

3.1 Dataset

In this section, we use TREC-QA dataset to evaluate the proposed model, which appears to be one of the most widely used benchmarks for answer sentence selection. This dataset was created by (Wang et al., 2007) based on Text REtrieval Conference (TREC) QA track (8-13) data[1]. Candidate answers were automatically retrieved for each factoid question. Two sets of data are provided for training, one is small training set containing 94 questions collected through manual judgement, and the other is full training set, i.e., Train-All, which is provided by (Wang et al., 2007) and contains 1,229 questions from the entire TREC 8-12 collection with automatically labeled ground truth by matching answer keys' regular expressions[2]. Table 1 summarizes the answer selection dataset in details. In the following experiments, we use the full training set due to its relatively large scale, even though there exist noisy labels caused by automatically pattern matching.

The original development and test datasets have 82 and 100 questions, respectively. Following (Wang and Nyberg, 2015; Santos et al., 2016; Tan et al., 2015), all questions with only positive or negative answers are removed. Finally, we have 65 development questions with 1,117 question-answer pairs, and 68 test questions with 1,442 question-answer pairs.

3.2 Token Representation

We use a pre-trained 50-dimensional word embeddings[3] (Pennington et al., 2014) as our initial word look-up table. These word embeddings are trained on Wikipedia data and the fifth English Gigawords with totally 6 Billion tokens. Since the proposed model is based on pairwise sentence similarity computing, 300-dimensional word embedding adopted by previous works (Wang and Nyberg, 2015; Santos et al., 2016; Tan et al., 2015) is too heavy for ours, which were trained by word2vec (Mikolov et al., 2013) on part of Google News dataset containing 100 billion words[4].

[1] http://cs.stanford.edu/people/mengqiu/data/qg-emnlp07-data.tgz

[2] http://cs.jhu.edu/~xuchen/packages/jacana-qa-naacl2013-data-results.tar.bz2

[3] http://nlp.stanford.edu/data/glove.6B.zip

[4] https://drive.google.com/file/d/0B7XkCwpI5KDYNlNUTTlSS21pQmM/

NetConfig	Choice
Embedding dim	50
Multi-Modal k	2 or 4
filter width in Layer 0	5
num. of filters in Layer 0	32
filter width in Layer 1	5
num. of filters in Layer 1	64
num. of hidden units	32
Dropout rate	0.5
Batch Norm	w

Table 2: Experimental choices of our net configuration. Layer represents for convolutional layer.

Model	MAP	MRR
M^{2}S-Net-Euc	.7265	.8218
M^{2}S-Net-Cos	.7413	.8322
M^{2}S-Net-Shallow	.7476	.8375
M^{2}S-Net-2	.7698	**.8640**
M^{2}S-Net-4	**.7793**	.8487

Table 3: Results of M^{2}S-Nets on the answer sentence selection dataset.

3.3 Experimental Setting

Following previous works, we also use the two metrics to evaluate the proposed model, i.e., Mean Average Precision (MAP) and Mean Reciprocal Rank (MRR). The official *trec_eval* scorer tool[5] is used to compute the above metrics.

We compute the simplest word overlap measurement between each question-answer pair, and concatenate it with our learned matching representation for the final rank learning. This feature vector contains only two features, i.e., word overlap and IDF-weighted word overlap.

We experiment our M^{2}S-Net on three predefined similarity measurements: M^{2}S-Net-Euc, M^{2}S-Net-Cos, and M^{2}S-Net-Metric. All of these three nets share the same network configuration (more details can be found in Table 2). To demonstrate the fact that the proposed network can benefit more from deep structure, we compare M^{2}S-Net-Metric with a one-convolutional layer network, namely M^{2}S-Net-Shallow. Furthermore, to verify the effectiveness of the proposed multi-modal similarity metric learning, we list results of M^{2}S-Net with $k = 2, 4$, respectively denoted

[5] http://trec.nist.gov/trec_eval/

Reference	MAP	MRR
Wang et al. (2007)	.6029	.6852
Heilman and Smith (2010)	.6091	.6917
Wang and Manning (2010)	.5951	.6951
Yao et al. (2013)	.6307	.7477
Yih et al. (2013)-LCLR	.7092	.7700
Yu et al. (2014)	.7113	.7846
Wang and Nyberg (2015)	.7134	.7913
Tan et al. (2015)	.7106	.7998
Severyn and Moschitti (2015)	.7459	.8078
Santos et al. (2016)	.7530	.8511
Wang et al. (2016)	.7714	.8447
M^2S-Net-2	.7698	**.8640**
M^2S-Net-4	**.7793**	.8487

Table 4: Results of our models and other methods from the literature.

as M^2S-Net-2 and M^2S-Net-4. All the networks mentioned here are implemented using Caffe (Jia et al., 2014).

4 Results and Discussion

From Table 3, we can see that the learned similarity metric works much better than the non-parametric hand-crafted similarity measurements. The results of shallow and deep network structure indicate that the proposed M^2S-Net benefits much from deep learning.

For comprehensive comparison, we also list the results of prior state-of-the-art methods in literature on this task in Table 4. It can be seen that the proposed method outperforms the most recently published attention-based methods by 1% in both MAP and MRR metrics.

5 Conclusion

We propose a novel end-to-end learning framework (M^2S-Net) for answer sentence selection task. Compared with the previous methods, which used artificial neural networks only to capture semantic representation of each sentence separately, we explore more on the interdependence between each sentence by constituting deep convolutional neural network based on multi-modal similarity metric learning on pairwise tokens. The proposed architecture is proved effective, and surpasses previous state-of-the-art systems on the answer selection benchmark, i.e., TREC-QA dataset, in both MAP and MRR metrics.

Acknowledgments

Thanks to Hanjie Wang for his good suggestions. Thanks to Kai Chen and Peng Shu for arrangement planing. Thanks to all the co-workers in Wireless Adtech RD team for advise and discussion.

References

[Bahdanau et al.2014] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

[Bitvai and Cohn2015] Zsolt Bitvai and Trevor Cohn. 2015. Non-linear text regression with a deep convolutional neural network. *Volume 2: Short Papers*, 1(x1):180.

[Denil et al.2014] Misha Denil, Alban Demiraj, Nal Kalchbrenner, Phil Blunsom, and Nando de Freitas. 2014. Modelling, visualising and summarising documents with a single convolutional neural network. *CoRR*, abs/1406.3830.

[Feng et al.2015] Minwei Feng, Bing Xiang, Michael R Glass, Lidan Wang, and Bowen Zhou. 2015. Applying deep learning to answer selection: A study and an open task. *arXiv preprint arXiv:1508.01585*.

[Heilman and Smith2010] Michael Heilman and Noah A Smith. 2010. Tree edit models for recognizing textual entailments, paraphrases, and answers to questions. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 1011–1019. Association for Computational Linguistics.

[Hu et al.2014] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *Advances in Neural Information Processing Systems*, pages 2042–2050.

[Ioffe and Szegedy2015] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.

[Jia et al.2014] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*.

[Kalchbrenner et al.2014] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. 2014. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, ACL 2014, June 22-27, 2014, Baltimore, MD, USA, Volume 1: Long Papers*, pages 655–665.

[Kim2014] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *EMNLP*.

[Mikolov et al.2013] Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. 2013. Linguistic regularities in continuous space word representations. In *HLT-NAACL*, pages 746–751.

[Pennington et al.2014] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

[Santos et al.2016] Cicero dos Santos, Ming Tan, Bing Xiang, and Bowen Zhou. 2016. Attentive pooling networks. *arXiv preprint arXiv:1602.03609*.

[Severyn and Moschitti2015] Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 373–382. ACM.

[Srivastava et al.2014] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958.

[Tan et al.2015] Ming Tan, Bing Xiang, and Bowen Zhou. 2015. Lstm-based deep learning models for non-factoid answer selection. *arXiv preprint arXiv:1511.04108*.

[Wang and Manning2010] Mengqiu Wang and Christopher D Manning. 2010. Probabilistic tree-edit models with structured latent variables for textual entailment and question answering. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 1164–1172. Association for Computational Linguistics.

[Wang and Nyberg2015] Di Wang and Eric Nyberg. 2015. A long short-term memory model for answer sentence selection in question answering. *ACL, July*.

[Wang et al.2007] Mengqiu Wang, Noah A Smith, and Teruko Mitamura. 2007. What is the jeopardy model? a quasi-synchronous grammar for qa. In *EMNLP-CoNLL*, volume 7, pages 22–32.

[Wang et al.2016] Zhiguo Wang, Haitao Mi, and Abraham Ittycheriah. 2016. Sentence similarity learning by lexical decomposition and composition. *arXiv preprint arXiv:1602.07019*.

[Yao et al.2013] Xuchen Yao, Benjamin Van Durme, Chris Callison-Burch, and Peter Clark. 2013. Answer extraction as sequence tagging with tree edit distance. In *HLT-NAACL*, pages 858–867. Citeseer.

[Yih et al.2013] Wen-tau Yih, Ming-Wei Chang, Christopher Meek, and Andrzej Pastusiak. 2013. Question answering using enhanced lexical semantic models. *ACL*.

[Yin et al.2015] Wenpeng Yin, Hinrich Schütze, Bing Xiang, and Bowen Zhou. 2015. Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *arXiv preprint arXiv:1512.05193*.

[Yu et al.2014] Lei Yu, Karl Moritz Hermann, Phil Blunsom, and Stephen Pulman. 2014. Deep learning for answer sentence selection. *NIPS Deep Learning and Representation Learning Workshop*.

[Zeiler and Fergus2013] Matthew D. Zeiler and Rob Fergus. 2013. Stochastic pooling for regularization of deep convolutional neural networks. *CoRR*, abs/1301.3557.

[Zeiler2012] Matthew D Zeiler. 2012. Adadelta: An adaptive learning rate method. *arXiv preprint arXiv:1212.5701*.

[Zeng et al.2014] Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. 2014. Relation classification via convolutional deep neural network. In *Proceedings of COLING*, pages 2335–2344.